

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
5 December 2002 (05.12.2002)

PCT

(10) International Publication Number
WO 02/098165 A1(51) International Patent Classification⁷: H04Q 11/04,
H04L 12/56

(21) International Application Number: PCT/SE02/01052

(22) International Filing Date: 3 June 2002 (03.06.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/870,945 1 June 2001 (01.06.2001) US

(71) Applicant: TELEFONAKTIEBOLAGET L M ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE).

(72) Inventors: KARLANDER, Bo; Idunvägen 5, S-191 43 Sollentuna (SE). AGNEVIK, Mikael; Oxenstierns väg 17, S-143 40 Gustavsberg (SE). WOLLBRAND, Per; Bergsundsgatan 13, S-117 37 Stockholm (SE).

(74) Agent: MAGNUSSON, Monica; Ericsson AB, Patent Unit Radio Access, S-164 80 Stockholm (SE).


(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: BANDWIDTH EFFICIENT QUALITY OF SERVICE SEPARATION OF AAL2 TRAFFIC



(57) Abstract: An AAL2 path group (60) comprises plural AAL2 paths (35). Bandwidth of an individual AAL2 path comprising the AAL2 path group is contributed to a total bandwidth of the AAL2 path group rather than to the individual AAL2 path exclusively. An admission decision regarding a connection seeking to use an AAL2 path belonging to the AAL2 path group is based on available bandwidth of the AAL2 path group rather than available bandwidth of an individual AAL2 path. ATM VCCs which comprise the AAL2 path group are transported on a virtual path (VP) together with ATM VCCs of a different type (e.g., a second type which differs from a first type of ATM VCC which comprise the AAL2 path group). The AAL2 path group also features quality of service (QoS) separation, e.g., differing treatment for differing AAL2 connections within the AAL2 path group based on the QoS requirements for the differing AAL2 connections.

BANDWIDTH EFFICIENT QUALITY OF SERVICE SEPARATION OF AAL2 TRAFFIC

BACKGROUND

1. FIELD OF THE INVENTION

5 The present invention pertains to wireless telecommunications, and particularly to transmission of AAL2 cells having differing Quality of Service (QoS) requirements.

2. RELATED ART AND OTHER CONSIDERATIONS

Aynchronous Transfer Mode (ATM) is now commonly used in communication networks. ATM is a packet-oriented transfer mode which uses asynchronous time
10 division multiplexing techniques. Packets are called cells and have a fixed size.

 An ATM cell consists of 53 octets, five of which form a header and forty eight of which constitute a "payload" or information portion of the cell. The header of the ATM cell includes two quantities which are used to identify a connection in an ATM network over which the cell is to travel, particularly the VPI (Virtual Path Identifier)
15 and VCI (Virtual Channel Identifier). In general, the virtual path is a principal path defined between two switching nodes of the network; the virtual channel is one specific connection on the respective principal path.

 Between termination points of an ATM network a plurality of nodes are typically situated, such as switching nodes having ports which are connected together
20 by physical transmission paths or links. The switching nodes each typically have several functional parts, a primary of which is a switch core. The switch core essentially functions like a cross-connect between ports of the switch. Paths internal to the switch core are selectively controlled so that particular ports of the switch are connected together to allow a cell ultimately to travel from an ingress side of the switch
25 to an egress side of the switch.

A protocol reference model has been developed for illustrating layering of ATM. The protocol reference model layers include (from lower to higher layers) a physical layer (including both a physical medium sublayer and a transmission convergence sublayer), an ATM layer, and an ATM adaptation layer (AAL), and higher layers. The basic purpose of the AAL layer is to isolate the higher layers from specific characteristics of the ATM layer by mapping the higher-layer protocol data units (PDU) into the information field of the ATM cell and vice versa. There are several differing AAL types or categories, including AAL0, AAL1, AAL2, AAL3/4, and AAL5.

AAL2 is a standard defined by ITU recommendation I.363.2. An AAL2 packet comprises a three octet packet header, as well as a packet payload. The AAL2 packet header includes an eight bit channel identifier (CID), a six bit length indicator (LI), a five bit User-to-User indicator (UUI), and five bits of header error control (HEC). The AAL2 packet payload, which carries user data, can vary from one to forty-five octets. Plural AAL2 packets can be inserted into a standard ATM cell. Thus, ATM Adaptation Layer Type 2, i.e., AAL2, facilitates multiplexing of plural AAL2 user connections on a common AAL2 path, established as a single ATM Virtual Channel Connection (VCC).

In an ATM based telecommunications system where different quality of service classes are supported, some connections are more delay sensitive than others. In order to cater to these differing sensitivities, ATM cells of differing priority must be handled. Such ATM cell handling can include specific traffic management per ATM-VCC with weighted fair queuing, early packet discard, available bit rate (ABR) accommodation, and shaping of outgoing traffic according to a traffic contract, for example. Handling of ATM traffic management is specified in the ITU I.371 Recommendation or in ATM-FORUM Traffic Management Specification 4.0.

The Universal Mobile Telecommunications (UMTS) Terrestrial Radio Access Network (UTRAN) is a third generation radio access network which, in some respects, builds upon the radio access technology known as Global System for Mobile communications (GSM) developed in Europe. The UTRAN covers a geographical area which is divided into cell areas, with each cell area being served by a base station. A cell is a geographical area where radio coverage is provided by the radio base station equipment at a base station site. Each cell is identified by a unique identity, which is broadcast in the cell. The base stations communicate over the air interface (e.g., radio

frequencies) with the user equipment units (UEs) within range of the base stations. In the radio access network, several base stations are typically connected (e.g., by landlines or microwave) to a radio network controller (RNC). The radio network controller, also sometimes termed a base station controller (BSC), supervises and coordinates various activities of the plural base stations connected thereto. The radio network controllers are typically connected to one or more core networks. UTRAN is essentially a wideband code division multiple access (W-CDMA) system.

In the UTRAN network, several types of control and user data connections use AAL2. These differing types of connections have differing quality of service (QoS) requirements, such as differing maximum delay tolerance, for example. However, standardized AAL2 does not provide any means for quality of service differentiation between connections on the same AAL2 path. This means that, when connections of different types share an AAL2 path, that shared path has to be dimensioned according to the quality of service requirements of the connection type having the most stringent such requirement (e.g., most stringent QoS).

In view of the foregoing, in some cases it is necessary to establish separate AAL2 paths for the more delay-sensitive traffic (such as conversational traffic, e.g., speech) and less delay-sensitive traffic (such as data traffic) in order to ensure that the most stringent requirements for the more delay-sensitive traffic can be met. In such cases, the AAL2 connections are typically divided into two quality of service (QoS) categories, and an AAL2 path is established to serve each of the AAL2 connection categories. Even the less delay-sensitive connections (e.g., data traffic connections) have strict maximum delay tolerance, although such tolerance is higher than for the delay-sensitive traffic (e.g., speech traffic).

In an implementation of the separate AAL2 paths as described in the preceding paragraph, one AAL2 path can, depending on the eight bit size of the connection identifier (CID), support a maximum of 248 AAL2 connections (some CID values are reserved for other purposes). When a higher number of connections must be supported between adjacent nodes, multiple AAL2 paths must be established.

Thus, conventional practice when handling AAL2 connections with differing QoS categories is to separate the different services on different AAL2 paths (ATM

VCCs) with reserved bandwidth on each AAL2 path. The ITU-T AAL2 standards include basic signaling methods to support such separation of AAL2 connections on different AAL2 path types and to set up these AAL2 paths on ATM VCCs with different ATM traffic contracts. For example, AAL2 cells carrying conversational (e.g., speech) traffic may be assigned to AAL2 paths with a first QoS class; AAL2 cells carrying non-conversational (e.g. data) traffic may be assigned to AAL2 paths with a second QoS class (lower than the first QoS class). AAL2 paths of both QoS classes are carried by the same ATM physical link.

In general, the type of ATM service category appropriate for AAL2 paths is Deterministic Bit Rate (DBR), which in the parlance of the ATM Forum is Constant Bit Rate (CBR). For this service category, ATM link resources are reserved according to the peak cell rate (PCR) of each VCC.

In a case in which several AAL2 paths are needed for QoS separation or connection capacity reasons, it is inefficient to allocate a fixed part of the link capacity to each AAL2 path. While the maximum total traffic over a physical link may be possible to estimate based on the capacity of the served radio interfaces (e.g., in the UTRAN), the relative proportion of conversational versus data traffic is more difficult to predict and is expected to change with time. If each path is dimensioned separately according to its expected maximum traffic intensity, more link capacity must be reserved than if the link were dimensioned according to the sum of the traffic on the AAL2 paths of the both types.

Saito, "Effectiveness of UBR VC Approach in AAL2 Networks and Its Application to IMT-2000", *IEICE Trans. Commun.*, Vol. E83-B, No. 11, November 2000, pp. 2486-2493, proposes a bandwidth management alternative which performs bandwidth management (using unspecified bit rate (UBR) for each VC) at a VP level (virtual path) rather than at a VC level. However, the Saito proposal has various limitations. For example, it assumes that the VP carries nothing other than AAL2 traffic, and also that all AAL2 connections are of the same QoS class. As an example, the Saito proposal does not take into consideration that other types of ATM VCCs, for example AAL5 connections carrying signaling, or operation and maintenance traffic could be included on the VP. Moreover, Saito does not address how differing QoS requirements for differing AAL2 connections should be handled.

What is needed, therefore, and an object of the present invention, is a technique which facilitates bandwidth efficient quality of service (QoS) separation of AAL2 traffic.

BRIEF SUMMARY OF THE INVENTION

5 The present invention concerns an AAL2 path group which comprises plural AAL2 paths with the same or different QoS classes. Bandwidth on the ATM layer is reserved for the AAL2 path group rather than to the individual AAL2 path exclusively. An admission decision regarding an AAL2 connection seeking to use an AAL2 path belonging to the AAL2 path group is based on available bandwidth of the AAL2 path
10 group rather than available bandwidth of the individual AAL2 path. ATM VCCs which comprise the AAL2 path group are transported on a virtual path (VP) together with ATM VCCs of a different type (e.g., a second type which differs from a first type of ATM VCCs which comprise the AAL2 path group).

15 The present invention also concerns an Asynchronous Transfer Mode (ATM) network which implement the AAL2 path group of the invention, as well as a node of such network. The Asynchronous Transfer Mode (ATM) network which implements the AAL2 path group comprises a first network node; a second network node; and, an ATM virtual path (VP) connecting the first network node and the second network node. One of the network nodes comprises a plurality of AAL2 queues; a path group
20 scheduler; non-path group ATM virtual circuits (VCs); and, a VP scheduler which applies ATM cells of the path group and the non-path group ATM virtual circuits (VCs) to the ATM virtual path (VP) for transmission to the second network node. Each of the plurality of AAL2 queues contains ATM cells from at least one AAL2 path. The path group scheduler selects ATM cells from the plurality of AAL2 queues to form the
25 AAL2 path group. The non-path group ATM virtual circuits (VCs) include ATM cells of a second type (e.g. AAL5) other than the ATM cells of a first type which comprise the AAL2 path group.

30 The AAL2 path group also features quality of service (QoS) separation, e.g., differing treatment for differing AAL2 connections within the AAL2 path group based on the QoS requirements for the differing AAL2 connections.

As one aspect of the invention, a connection admission controller unit admits connections to the AAL2 paths comprising the AAL2 path group based on available bandwidth of the AAL2 path group rather than available bandwidth of an individual AAL2 path. The connection admission controller unit can be situated at the same network node as the plurality of AAL2 queues; the path group scheduler; the non-path group queue; and, the VP scheduler. Alternatively, the connection admission controller unit can be situated at a node other than the network node which has the plurality of AAL2 queues; the path group scheduler; the non-path group queue; and, the VP scheduler.

In the present invention, ATM cells of the first type comprising the AAL2 path group are transmitted over the same virtual path (VP) as ATM cells of a second, non-AAL2 type. Advantageously, the first type of ATM cells and second type of ATM cells can be delineated in various different ways. In one mode of the invention, the first type of ATM cells comprises AAL2 traffic cells and the second type of ATM cells comprises cells dedicated to for example signaling, operation and maintenance, or synchronization. The ATM cells of the second type thus can be cells which require a guaranteed bandwidth. In another mode the VP only comprises ATM cells of the first type. Additionally, the first type of ATM cells can comprise cells from AAL2 paths with different QoS classes. For example, the first type of ATM cells can comprise AAL2 cells having an unspecified bit rate (UBR) and the second type of ATM cells can comprise AAL2 cells have a bit rate type other than UBR.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the following more particular description of preferred embodiments as illustrated in the accompanying drawings in which reference characters refer to the same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Fig. 1 is a schematic view of a telecommunications system comprising nodes, with at least one of the nodes being partly illustrated in example fashion as forming an AAL2 path group in accordance with the present invention.

Fig. 1A is a schematic view of a variation of the telecommunications system of Fig. 1 wherein a node which forms an AAL2 path group also includes admission control.

Fig. 1B is a schematic view of a variation of the telecommunications system of Fig. 1 wherein a node which forms an AAL2 path group does not include admission control.

Fig. 2A - Fig. 2C are schematic view of portions of nodes which form an AAL2 path group in accordance with differing embodiments of the present invention.

Fig. 3 is a schematic view of a variation of the telecommunications system of Fig. 1 which employs a different pre-AAL2 path group scheduler front end.

Fig. 4 is a partially schematic, partially diagrammatic view showing an admission control scenario according to a mode of the invention in an example context of Fig. 1A.

DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well-known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail. Moreover, individual function blocks are shown in some of the figures. Those skilled in the art will appreciate that the functions may be implemented using individual hardware circuits, using software functioning in conjunction with a suitably programmed digital microprocessor or general purpose computer, using an application specific integrated circuit (ASIC), and/or using one or more digital signal processors (DSPs).

Fig. 1 shows, in representative, non-limiting fashion, a telecommunications system 20 having two example nodes 22A and 22B which are connected by an ATM virtual path (VP) realized by physical link 24. As described in more detail below,

telecommunications system 20 implements an AAL2 path group according to the present invention. In the illustrated embodiment, node 22A is broadly configured to implement the inventive AAL2 path group, and in particular with regard to incoming ATM virtual circuits (VCs) shows a first plurality of incoming AAL2 connections 30₁ (comprising AAL2 connections 30₁₋₁ through 30_{1-k} for carrying AAL2 traffic); a second plurality of incoming AAL2 connections 30₂ (comprising AAL2 connections 30₂₋₁ through 30_{2-j} also for carrying AAL2 traffic); and, an ATM VC 30_a which carries ATM cells of a different type than are carried on incoming AAL2 connections 30₁ and 30₂.

Node 22A also comprises an AAL2 multiplexer for each plurality of incoming AAL2 connections (e.g., AAL2 multiplexer 32₁ for plurality 30₁ and AAL2 multiplexer 32₂ for plurality 30₂); an AAL2 path 35 output which is output from each AAL2 multiplexer; a plurality of AAL2 queues represented by AAL2 queues 40₁ and 40₂; AAL2 path group scheduler 42; and VP scheduler 52. In the illustrated scenario of Fig. 1, inputs to queues 40₁ and 40₂ are connected to AAL2 paths 35₁ and 35₂ which, in turn, are fed by AAL2 multiplexer 32₁ and 32₂, respectively. AAL2 path group scheduler 42 selects cells from queues 40₁ and 40₂ in accordance with a predetermined selection strategy or logic. The AAL2 paths which feed the queues 40₁, 40₂ from which AAL2 path group scheduler 42 selects ATM cells comprise an AAL2 path group 60, and thus AAL2 cells output by AAL2 path group scheduler 42 on line 61 comprise cells of the AAL2 path group 60. The VP scheduler 52 applies ATM cells of the AAL2 path group 60 as well as cells from one or more non-path group VCs (e.g., ATM VC 30_a) to the ATM virtual path (VP) 24 for transmission to the second network node 22B.

Each of the plurality of AAL2 queues 40₁, 40₂, contains ATM AAL2 cells from at least one AAL2 path. The AAL2 path group scheduler 42 selects ATM AAL2 cells from the plurality of AAL2 queues 40₁, 40₂, with the AAL2 cells discharged by AAL2 path group scheduler 42 on line 61 constituting cells of the AAL2 path group. The bandwidth of an individual AAL2 path comprising the AAL2 path group 60 is contributed to a total bandwidth of the AAL2 path group rather than to the individual AAL2 path exclusively.

In one example context of implementation of the present invention the CID of an AAL2 cell stream is employed to identify a user traffic connection. In the present invention, an admission decision regarding a user seeking to establish a connection

which employs a certain one of the AAL2 paths belonging to the AAL2 path group is based on available bandwidth of the entire AAL2 path group 60, rather than available bandwidth of that certain AAL2 path individually. That admission decision is made by an admission control function, which has access to the capacities and capabilities of the network including the collective (e.g., total) bandwidth of the AAL2 path group (e.g., the total bandwidth of the AAL2 paths comprising the AAL2 path group).

In an embodiment depicted in Fig. 1A, the admission control function is shown as an admission controller 100A which is located in node 22A, i.e., in the same node with the structure (e.g., AAL2 path group scheduler 42) which forms the AAL2 path group sending cells out from node 22A. The admission controller 100A is connected to and accesses a database, referenced for simplicity herein as "capacity database" 102A. The capacity database 102A stores configuration parameters for network 20 and for various constituent nodes, including the bandwidth capacity for each of the AAL2 paths 35₁ and 35₂. As indicated by arrow 104A, admission controller 100A is configured to realize that the AAL2 paths 35₁ and 35₂ comprise the AAL2 path group 60 of the present invention, and thus any admission decision, such as a decision whether to admit a new connection over one of the AAL2 paths 35₁ and 35₂, is assessed with respect to the total available bandwidth of the AAL2 path group 60 rather than to the available bandwidth of any particular AAL2 path. Thus, it is not necessary for the capacity database to store a bandwidth for each individual AAL2 path.

In the embodiment of Fig. 1A, node 22A may be a radio network control (RNC) node such as is employed in the radio access network of the Universal Mobile Telecommunications (UMTS) Terrestrial Radio Access Network (UTRAN), as admission control functions are generally performed by RNC nodes in UTRAN. In such embodiment, node 22B can be any other node connected to a RNC, such as a base station (BS) node (e.g., Node B), for example. In the case of node 22B being a base station (BS) node, the base station (BS) node 22B may have structure comparable to node 22A for forming a AAL2 path group which is transmitted from node 22B to node 22A, with admission control of connections involving segments between node 22A and node 22B being controlled by either admission controller 100A or a local admission controller in node 22B.

In another embodiment depicted in Fig. 1B, the admission control function is shown as an admission controller 100B which is located in node 22B. In other words, node 22A with the structure (e.g., AAL2 path group scheduler 42) which forms the AAL2 path group sending cells does not have an admission control function in the same node. Rather, as illustrated by arrow 104B, admission control regarding the AAL2 path group comprised of AAL2 paths 35_1 and 35_2 is performed by admission controller 100B situated in node 22B. The node 22A of the embodiment of Fig. 1B can be, for example, a base station (BS) node of a UTRAN, as alluded to in the preceding paragraph.

Fig. 2A - Fig. 2C show, in accordance with differing embodiments, examples of how node 22A can be configured to form an AAL2 path group in accordance with the present invention. While the admission control aspect is not explicitly illustrated in Fig. 2A - Fig. 2C, it should be understood that admission control can be provided in various manners (such as inclusion of admission controller 100A in node 22A as shown in Fig. 1A or inclusion of admission controller 100B in node 22B as shown in Fig. 1B).

Node 22A of Fig. 2A has a AAL2 path group scheduler 42A which is fed in the same manner as Fig. 1, which means that the AAL2 path group comprises the AAL2 paths 35_1 and 35_2 . Fig. 2A (like Fig. 2B and Fig. 2C) specifies that the AAL2 path 35_1 is comprised of speech (e.g., conversational) AAL2 connections, while the AAL2 path 35_2 is comprised of data (e.g., non-speech or non-conversational) AAL2 connections. In the Fig. 2A embodiment, AAL2 path group scheduler 42A is a Weighted Fair Queue scheduler which selects cells from queue 40_1 (e.g., cells with speech payloads) according to a weight of X% and which selects cells from queue 40_2 (e.g., cells with data payloads) according to a weight of Y%. The circle symbol employed at the output port of AAL2 path group scheduler 42A is utilized in the industry to denote a traffic shaping function performed by AAL2 path group scheduler 42A. In particular, the outgoing traffic is shaped to the peak cell rate (PCR).

Whereas in Fig. 1 ATM cells of a different type than those comprising the AAL2 path group are generically depicted as being applied on path 30_a , Fig. 2A - Fig. 2C show with more specificity some examples of such differing types of ATM cells. For example, Fig. 2A shows VC 30_{OAM} which carries cells utilized for operation and maintenance purposes (e.g., operator commands, maintenance statistics, alarms, etc.).

Cells received on VC 30_{OAM} are applied to queue 40_{OAM}, with an output of queue 40_{OAM} being applied to intermediate scheduler 62. As another example, Fig. 2A shows an VC 30_{SIGNALING} which carries cells utilized for signaling purposes (e.g., for setting up and tearing down connections, including admission control signaling). Cells received on VC 30_{SIGNALING} are applied to queue 40_{SIGNALING}, with an output of queue 40_{SIGNALING} also being applied to intermediate scheduler 62. As yet another example, Fig. 2A shows an VC 30_{SYNC} which carries cells utilized for synchronization purposes. Cells received on VC 30_{SYNC} are applied to queue 40_{SYNC}, with an output of queue 40_{SYNC} being applied to VP scheduler 52.

The intermediate scheduler 62 thus receives, as input, cells forming AAL2 path group as output on line 61 from AAL2 path group scheduler 42A, as well as OAM cells from queue 40_{OAM} and cells with signaling information from queue 40_{SIGNALING}. In the illustrated embodiment, intermediate scheduler 62 uses Weighted Fair Queueing, and accords cells comprising the AAL2 path group a weight of 80%, the OAM cells a weight of 5%, and the signaling cells from queue 40_{SIGNALING} a weight of 15%. The illustrated percentage weightings are merely illustrative, as other percentage weightings can be employed in other embodiments. The output of intermediate scheduler 62 is applied to a first input of VP scheduler 52; synchronization cells stored in queue 40_{SYNC} are applied to a second input of VP scheduler 52. The VP scheduler 52 of the Fig. 2A embodiment is a strict priority input scheduler which schedules application of outgoing cells onto virtual path 24 in accordance with the direction of the arrow (e.g., with synchronization cells in queue 40_{SYNC} always having a highest priority).

Node 22B of the embodiment of Fig. 2B is similar to node 22A of the embodiment of Fig. 2A, but differs primarily in that AAL2 path group scheduler 42B is a strict priority input scheduler which always gives cells queue 40₁ (e.g., cells having speech payload) precedence or priority over cells in queue 40₂ (e.g., cells having data payload). The order of the priority is indicated by the downward pointing arrow symbolized in the interior of AAL2 path group scheduler 42B.

Node 22C of the embodiment of Fig. 2C is also similar to node 22A of the embodiment of Fig. 2A, but its scheduler 42C is not dedicated to scheduling the AAL2 path group only. Rather, scheduler 42C is a Fair Weight Queue scheduler which also consolidates therein the functions of intermediate scheduler 62. The output of AAL2

path group scheduler 42C is applied to a first input of VP scheduler 52; the synchronization cells stored in queue 40_{SYNC} are applied to a second input of VP scheduler 52. The VP scheduler 52 of the Fig. 2C embodiment is a strict priority input scheduler which schedules application of outgoing cells onto virtual path 24 in accordance with the direction of the arrow (e.g., with synchronization cells in queue 40_{SYNC} always having a highest priority).

The embodiments previously illustrated have shown, by way of example, a dedicated queue 40 for each AAL2 path. However, it should be understood that the particular elements comprising a front end leading to the input(s) of AAL2 path group scheduler 42 are not critical to the present invention and that other configurations are within the scope of the invention. For example, Fig. 3 illustrates another example embodiment in which plural AAL2 paths carrying AAL2 connections with the same quality of service (QoS) requirements are queued together. As shown in Fig. 3, queue 40₁ is fed by AAL2 paths 35₁₋₁ and 35₁₋₂, while queue 40₂ is fed by AAL2 paths 35₂₋₁ and 35₂₋₂. In the example illustration of Fig. 3, each AAL2 path carries plural AAL2 connections having the same quality of service (QoS) requirements. For example, AAL2 path 35₁₋₁ carries AAL2 connection 30₁₋₁₋₁ through 30_{1-1-k}, all of these connections having the same QoS. Moreover, AAL2 path 35₁₋₂ carries AAL2 connection 30₁₋₂₋₁ through 30_{1-2-k}, all of these connections having the same QoS.

Thus, there can be a number of connections/queues per traffic type. Common to all illustrated embodiments, however, is the fact that more delay sensitive traffic (e.g., conversational or speech traffic) and less delay sensitive traffic (e.g., non-conversational or data traffic) are separated into different queues to prevent head of line blocking.

In the present invention, ATM cells of the first type comprising the AAL2 path group 60 are transmitted over a same virtual path (VP) (e.g., virtual path 24) with ATM cells of a second type. Advantageously, the first type of ATM cells and second type of ATM cells can be delineated in various different ways.

In one mode of the invention, the first type of ATM cells which are carried over the paths comprising the AAL2 path group are AAL2 traffic cells while the second type of ATM cells comprises cells dedicated to at least one of signaling, operation and

maintenance, and synchronization. The embodiments of Fig. 2A - Fig. 2C typify such mode, since cells from VC 30_{OAM}, VC 30_{SIGNALING}, and VC 30_{SYNC} are all eligible for scheduling by VP scheduler 52, along with the cells carried by AAL2 path group, for application on virtual path (VP) 24. The ATM cells of the second type thus can be cells which require a guaranteed bandwidth.

As an alternative way of describing this phenomena, the first type of ATM cells can comprise AAL2 cells and the second type of ATM cells can comprise non-AAL2 cells (e.g., ATM cells that have another adaptation layer type). As a further alternative, the first type of ATM cells can comprise cells from AAL2 paths with differing QoS classes. For example, the first type of ATM cells which comprise the AAL2 path group can be AAL2 cells having an unspecified bit rate (UBR) and the second type of ATM cells can comprise AAL2 cells have a bit rate type other than UBR.

As mentioned above, in the present invention several AAL2 paths on the same physical link are handled together in an AAL2 path group with an assigned bandwidth. The bandwidth on the ATM layer is reserved for the AAL2 path group, rather than for the individual AAL2 path. Then, at establishment of an AAL2 connection, the connection admission control is calculated against the available bandwidth of the AAL2 path group and not against the bandwidth of the chosen AAL2 path. CID assignment within AAL2 is performed in the normal way, e.g., chosen for each individual AAL2 path. As illustrated in various foregoing embodiments, conversational (e.g., speech) traffic and non-conversational traffic (e.g., data) are separated on different AAL2 paths, and also queued separately.

Fig. 4 shows, in the context of the embodiment of Fig. 1A, admission of a new connection to the AAL2 path group in an illustrative scenario. The capacity database 102 shows, for each AAL2 path, a total capacity and an available capacity. For example, AAL2 path 35₁ has a total capacity TC₁ but no available capacity, whereas AAL2 path 35₂ has a total capacity TC₂ and an available capacity AC₂. At the time shown in Fig. 4, a new connection 30_{1,1} is sought over AAL2 path 35₁. Since AAL2 path 35₁ has no available capacity, traditional admission control techniques would either reject the requested new connection 30_{1,1} or cut off an existing connection carried on AAL2 path 35₁ in order to make room for the requested new connection. In accordance with the present invention, by contrast, admission controller 100 realizes

that, while AAL2 path 35₁ has no available capacity, the AAL2 path group 60 which includes AAL2 path 35₁ has available capacity. For example, AAL2 path 35₂ is one of the AAL2 paths comprising the AAL2 path group 60 which has spare capacity. Assuming the total spare capacity of the AAL2 path group exceeds the bandwidth
5 required for the requested new connection 30₁₋₁, in accordance with the present invention the admission controller 100 admits the new requested connection.

While the admission control scenario of Fig. 4 has been described in the context of the embodiment of Fig. 1A, it should be realized that the present invention is not limited by such context, and that other contexts such as the embodiment of Fig. 1B are
10 also within the purview of the present invention.

As reflected in the embodiments hereinabove described, the AAL2 path group also features quality of service (QoS) separation, e.g., differing treatment for differing AAL2 connections within the AAL2 path group based on the QoS requirements for the differing AAL2 connections. For example, AAL2 connections having differing QoS
15 requirements can be queued separately upstream of the AAL2 path group scheduler 42. Moreover, the AAL2 path group scheduler 42 can provide a preference or weighting in view of the QoS parameters for the respective AAL2 connections. In the embodiments of Fig. 2A and Fig. 2C, for example, the AAL2 path group scheduler 42A and AAL2 path group scheduler 42C, respectively, can weight the AAL2 VCs having speech
20 traffic payload more heavily than the AAL2 VCs having data traffic payloads. Or in the embodiment of Fig. 2B, the AAL2 path group scheduler 42B by its strict priority nature assigns a higher priority to the AAL2 VCs having speech traffic payload than the AAL2 VCs having data traffic payloads.

The present invention assigns bandwidth to its AAL2 path group instead of each
25 individual AAL2 path, thereby affording efficient bandwidth utilization. Allocation of bandwidth resources for AAL2 connections is thus made against the pooled resources of multiple AAL2 paths rather than against a fixed resource per path. In the present invention, the AAL2 paths themselves are treated as having UBR (unspecified bit rate), but the AAL2 path group is considered as a whole to have CBR (constant bit rate).
30 ATM layer transmission scheduling is also performed in such a way that all AAL2 paths are scheduled together as a single stream of ATM cells.

The present invention advantageously accommodates both UBR and non-UBR VCs (assigned to something other than AAL2) when performing AAL2 connection admission control against the bandwidth of the virtual path (VP). The invention does not require that only AAL2 path VCs be setup within a virtual path. By contrast, prior art techniques which perform connection admission control for AAL2 against the virtual path bandwidth must dedicate the entire virtual path to AAL2.

Moreover, the present invention only requires one virtual path between two nodes, which is economically desirable when transport is leased from an ATM transport network provider. Leasing one virtual path of, e.g., 2 Mbit/s is much cheaper than leasing two virtual paths with 1 Mbit/s each. Further, the present invention provides a better multiplexing gain on the ATM level when all cells are multiplexed on the same virtual path.

In one variation of the invention, different services can be separated on different virtual paths, although such variation may not yield efficient bandwidth utilization when using public VP networks, with coarse granularity of the possible rates.

The present invention is particularly applicable to AAL2 nodes, but advantageously does not have any impact on intermediate, pure ATM nodes, provided that such intermediate ATM nodes must have virtual path switches (e.g., cross-connects).

While embodiments of the present invention have already been said to be applicable to nodes such as radio network control (RNC) nodes and base station (BS) nodes, the invention is not limited to such nodes. For example, the present invention can also be utilized with other nodes such as media gateways.

While the invention has been described in connection with what is presently considered to be the most practical and preferred embodiment, it is to be understood that the invention is not to be limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.

WHAT IS CLAIMED IS:

- 1 1. A method of managing traffic for transport on an Asynchronous Transfer
2 Mode (ATM) virtual path (VP), the method characterized by:
3 forming an AAL2 path group which comprises plural AAL2 paths;
4 admitting connections based on available bandwidth of the AAL2 path group
5 rather than available bandwidth of an individual AAL2 path;
6 including, for transport on the virtual path (VP), ATM cells of a second type, the
7 cells of the second type being other than a first type of ATM cells, the first type of
8 ATM cell being cells which comprise the AAL2 path group.
- 1 2. The method of claim 1, wherein the first type of ATM cells comprises AAL2
2 traffic cells and the second type of ATM cells comprises cells dedicated to at least one
3 of signaling, operation and maintenance, and synchronization.
- 1 3. The method of claim 1, wherein the second type of ATM cells comprises
2 cells which require a guaranteed bandwidth.
- 1 4. The method of claim 1, wherein the first type of ATM cells comprises AAL2
2 cells and the second type of ATM cells comprises non-AAL2 cells.
- 1 5. The method of claim 1, wherein the first type of ATM cells comprises AAL2
2 cells from AAL2 paths with differing QoS classes.
- 1 6. The method of claim 5, wherein the first type of ATM cells comprises AAL2
2 cells having an unspecified bit rate (UBR) and the second type of ATM cells comprises
3 AAL2 cells have a bit rate type other than UBR.
- 1 7. The method of claim 1, wherein a given one of the AAL2 paths carries
2 AAL2 cells having a same Quality of Service (QoS) requirement.
- 1 8. The method of claim 1, further comprising allocating more delay-sensitive
2 traffic to a different AAL2 path than less delay-sensitive traffic.

1 9. The method of claim 1, further comprising providing differing treatment for
2 differing AAL2 connections within the AAL2 path group based on the QoS
3 requirements for the differing AAL2 connections.

1 10. The method of claim 1, further comprising providing a greater weighting or
2 priority for more delay sensitive AAL2 connections.

1 11. A node (22) of an Asynchronous Transfer Mode (ATM) characterized by:
2 a plurality of AAL2 paths (35) which together form an AAL2 path group (60);
3 and

4 an admission controller (100) arranged to admit a connection to one of the
5 AAL2 paths of the AAL2 path group based on available bandwidth of the AAL2 path
6 group rather than available bandwidth of an individual AAL2 path.

1 12. The apparatus of claim 11, further comprising
2 a path group scheduler (42) arranged to select ATM AAL2 cells from the
3 plurality of AAL2 paths;
4 a non-path group ATM virtual circuit (VC) arranged to include ATM cells of a
5 second type, the cells of the second type being other than ATM cells of a first type, the
6 ATM cells of the first type comprising the AAL2 path group;
7 a scheduler (52) arranged to apply ATM cells of the path group and the non-
8 path group ATM virtual circuit (VC) to an ATM virtual path (VP) for transmission to
9 away from the node.

1 13. The apparatus of claim 12, wherein the first type of ATM cells comprises
2 AAL2 traffic cells and the second type of ATM cells comprises cells dedicated to at
3 least one of signaling, operation and maintenance, and synchronization.

1 14. The apparatus of claim 12, wherein the second type of ATM cells
2 comprises cells which require a guaranteed bandwidth.

1 15. The apparatus of claim 12, wherein the first type of ATM cells comprises
2 AAL2 cells and the second type of ATM cells comprises non-AAL2 cells.

1 16. The apparatus of claim 12, wherein the first type of ATM cells comprises
2 AAL2 cells from AAL2 paths with differing QoS classes.

1 17. The apparatus of claim 16, wherein the first type of ATM cells comprises
2 AAL2 cells having an unspecified bit rate (UBR) and the second type of ATM cells
3 comprises AAL2 cells have a bit rate type other than UBR.

1 18. The apparatus of claim 12, wherein a given one of the AAL2 paths carries
2 AAL2 cells having a same Quality of Service (QoS) requirement.

1 19. The apparatus of claim 11, further comprising allocating more delay-
2 sensitive traffic to a different AAL2 path than less delay-sensitive traffic.

1 20. The apparatus of claim 12, wherein the path group scheduler is a weighted
2 fair queuing scheduler.

1 21. The apparatus of claim 12, wherein the path group scheduler is a strict
2 priority scheduler.

1 22. The apparatus of claim 11, further comprising means for providing
2 differing treatment for differing AAL2 connections within the AAL2 path group based
3 on the QoS requirements for the differing AAL2 connections.

1 23. The apparatus of claim 22, further comprising means for providing a greater
2 weighting or priority for more delay sensitive AAL2 connections.

1 24. An Asynchronous Transfer Mode (ATM) network (20) comprising a first
2 network node (22A) and a second network node (22B) with an ATM virtual path (VP)
3 connecting the first network node and the second network node;

4 characterized in that the first network node comprises:
5 a plurality of AAL2 paths (35) which together form an AAL2 path group
6 (60) wherein a connection is admitted to one of the AAL2 paths of the AAL2 path
7 group is based on available bandwidth of the AAL2 path group rather than available
8 bandwidth of an individual AAL2 path;

9 a path group scheduler (42) arranged to select ATM AAL2 cells from the
10 plurality of AAL2 paths;
11 a non-path group ATM virtual circuit (VC) arranged to carry ATM cells
12 of a second type, the ATM cells of the second type being other than ATM cells of a
13 first type, the ATM cells of the first type comprising the AAL2 path group;
14 a scheduler (52) arranged to apply ATM cells of the path group and the
15 non-path group ATM virtual circuit (VC) to the ATM virtual path (VP) for
16 transmission to the second network node.

1 25. The apparatus of claim 24, further comprising a connection admission
2 controller unit (100) arranged to admit connections based on available bandwidth of the
3 AAL2 path group rather than available bandwidth of an individual AAL2 path.

1 26. The apparatus of claim 25, wherein the connection admission controller unit
2 is situated at the first network node (22A).

1 27. The apparatus of claim 25, wherein the connection admission controller unit
2 is situated at the second network node (22B).

1 28. The apparatus of claim 24, wherein the first type of ATM cells comprises
2 AAL2 traffic cells and the second type of ATM cells comprises cells dedicated to at
3 least one of signaling, operation and maintenance, and synchronization.

1 29. The apparatus of claim 24, wherein the second type of ATM cells
2 comprises cells which require a guaranteed bandwidth.

1 30. The apparatus of claim 24, wherein the first type of ATM cells comprises
2 AAL2 cells and the second type of ATM cells comprises non-AAL2 cells.

1 31. The apparatus of claim 24, wherein the first type of ATM cells comprises
2 AAL2 cells having a first bit rate type and the second type of ATM cells comprises
3 AAL2 cells having a second bit rate type.

1 32. The apparatus of claim 31, wherein the first type of ATM cells comprises
2 AAL2 from AAL2 paths with differing QoS classes.

1 33. The apparatus of claim 24, wherein a given one of the AAL2 paths carries
2 AAL2 cells having a same Quality of Service (QoS) requirement.

1 34. The apparatus of claim 24, further comprising allocating more delay-
2 sensitive traffic to a different AAL2 path than less delay-sensitive traffic.

1 35. The apparatus of claim 24, wherein one of the first network node and the
2 second network node is a base station node.

1 36. The apparatus of claim 24, wherein the one of the first network node and the
2 second network node is a radio network controller node.

1 37. The apparatus of claim 24, wherein the path group scheduler is a weighted
2 fair queuing scheduler.

1 38. The apparatus of claim 24, wherein the path group scheduler is a strict
2 priority scheduler.

1 39. The apparatus of claim 24, further comprising means for providing
2 differing treatment for differing AAL2 connections within the AAL2 path group based
3 on the QoS requirements for the differing AAL2 connections.

1 40. The apparatus of claim 39, further comprising means for providing a greater
2 weighting or priority for more delay sensitive AAL2 connections.

1 41. Apparatus for managing traffic for transport on an Asynchronous Transfer
2 Mode (ATM) virtual path (VP), the apparatus characterized by:

3 means for forming an AAL2 path group (60) which comprises plural AAL2
4 paths (35);

5 means (100) for admitting connections based on available bandwidth of the
6 AAL2 path group rather than available bandwidth of an individual AAL2 path;

7 means (52) for including, for transport on the virtual path (VP), ATM cells of a
8 second type, the ATM cells of the second type being other than a first type of ATM
9 cells, the first type of ATM cells comprising the AAL2 path group.

1 42. The apparatus of claim 41, wherein the first type of ATM cells comprises
2 AAL2 traffic cells and the second type of ATM cells comprises cells dedicated to at
3 least one of signaling, operation and maintenance, and synchronization.

1 43. The apparatus of claim 41, wherein the second type of ATM cells
2 comprises cells which require a guaranteed bandwidth.

1 44. The apparatus of claim 41, wherein the first type of ATM cells comprises
2 AAL2 cells and the second type of ATM cells comprises non-AAL2 cells.

1 45. The apparatus of claim 41, wherein the first type of ATM cells comprises
2 AAL2 cells from AAL2 paths with differing QoS classes.

1 46. The apparatus of claim 45, wherein the first type of ATM cells comprises
2 AAL2 cells having an unspecified bit rate (UBR) and the second type of ATM cells
3 comprises AAL2 cells have a bit rate type other than UBR.

1 47. The apparatus of claim 41, wherein a given one of the AAL2 paths carries
2 AAL2 cells having a same Quality of Service (QoS) requirement.


1 48. The apparatus of claim 41, further comprising allocating more delay-
2 sensitive traffic to a different AAL2 path than less delay-sensitive traffic.

1 49. The apparatus of claim 41, further comprising means for providing
2 differing treatment for differing AAL2 connections within the AAL2 path group based
3 on the QoS requirements for the differing AAL2 connections.

1 50. The apparatus of claim 49, further comprising means for providing a greater
2 weighting or priority for more delay sensitive AAL2 connections.

1/8

Fig. 1



2/8

Fig. 1A




Fig. 1B

3/8







Fig. 2A



6/8


Fig. 2C





8/8

Fig. 4



INTERNATIONAL SEARCH REPORT

 Int. Application No.
 PCT/SE 02/01052

 A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04Q11/04 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EP0-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	HIROSHI SAITO: "Bandwidth management for AAL2 traffic." IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, vol. 49, no. 4, July 2000 (2000-07), pages 1364-1377, XP002902637 ISSN: 0018-9545 page 1366, paragraph III -page 1367 ---	1-50
P,A	US 2001/030966 A1 (CHOI MYUNG SOON) 18 October 2001 (2001-10-18) column 1, line 52 -column 2, line 32 abstract ---	1-50
A	WO 98 27692 A (MOTOROLA INC) 25 June 1998 (1998-06-25) page 15, line 25 -page 17, line 19 abstract --- -/-	1-50

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
 "E" earlier document but published on or after the international filing date
 "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
 "O" document referring to an oral disclosure, use, exhibition or other means
 "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
 "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
 "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
 "Z" document member of the same patent family

Date of the actual completion of the international search

27 August 2002

Date of mailing of the international search report

11 09. 2002

Name and mailing address of the ISA

 European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx 31 051 epo nl,
 Fax: (+31-70) 340-3010

Authorized officer

Peter Hedman

INTERNATIONAL SEARCH REPORT

Int. Patent Application No.
PCT/SE 02/01052

C. (Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 915 634 A (LUCENT TECHNOLOGIES INC) 12 May 1999 (1999-05-12) column 10, line 20 -column 14, line 23 abstract -----	1-50

INTERNATIONAL SEARCH REPORT

Information on patent family members

Initial Application No

PCT/SE 02/01052

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 2001030966	A1	18-10-2001	NONE	
WO 9827692	A	25-06-1998	AU 5167498 A CN 1240081 A EP 0947075 A1 WO 9827692 A1	15-07-1998 29-12-1999 06-10-1999 25-06-1998
EP 0915634	A	12-05-1999	US 6169738 B1 EP 0915634 A2 JP 11225150 A	02-01-2001 12-05-1999 17-08-1999